

Automating Coordinated Autonomous Vehicle Control

Extended Abstract

Allen Huang

Department of Computer Science
University of Cape Town, Cape Town, South Africa
allen@allenuang.net

Geoff Nitschke

Department of Computer Science
University of Cape Town, Cape Town, South Africa
gnitschke@cs.uct.ac.za

ACM Reference Format:

Allen Huang and Geoff Nitschke. 2020. Automating Coordinated Autonomous Vehicle Control. In *Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), Auckland, New Zealand, May 9–13, 2020*, IFAAMAS, 2 pages.

1 INTRODUCTION

Recently there has been increasing academic and industry research attention on producing adaptive control systems for autonomous vehicles. To accommodate such autonomous vehicles there have been proposals that current road and highway infrastructure undergo significant changes. For example, replacing traffic lights and stop signs and allowing autonomous vehicles to coordinate their own interactions so as to avoid collisions and safely navigate through intersections [1]. One approach is to design vehicle controllers such that desired multi-agent behaviors (coordinated driving behaviors) are automatically synthesized for vehicles driving on any given road network [3].

This study investigates evolutionary controller design for enabling effective and efficient coordinated driving behavior for autonomous vehicle traffic operating on roads built exclusively for autonomous vehicles. That is, roads and highways without the current road infrastructure of traffic lights, intersection stop signals and vehicle lanes [1]. Vehicle controllers must coordinate their driving behaviors so as all vehicle traffic *effectively* and *efficiently* passes through increasingly difficult road networks. Effectiveness and efficiency are vehicle traffic task performance metrics, equating to the number of collisions and time taken to traverse a given road, and must thus be minimized. Task difficulty is the number of vehicles (traffic density) and obstacles on the road. Specifically, a goal of this work is to automate the synthesis of vehicle controllers such that when multiple vehicles interact a desired coordinated driving behavior emerges for any given task environment (road, vehicles and obstacles).

This study’s main contribution is the evolutionary synthesis of coordinated, effective and efficient vehicle traffic. First, *objective-based search*, second, *non-objective based search*, and third *Hybrid search* combining the objective and non-objective-based search functions. Vehicle controller evolution was the coupling of an objective, non-objective or hybrid evolutionary search method with NEAT [6] for adapting neural controllers.

Proc. of the 19th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2020), B. An, N. Yorke-Smith, A. El Fallah Seghrouchni, G. Sukthankar (eds.), May 9–13, 2020, Auckland, New Zealand. © 2020 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). All rights reserved.

2 METHODS

Vehicle controllers were evolved with one of three NEAT variants (*fitness function*, *novelty search*, *hybrid search*), with the goal of maximizing the average distance traversed on a given *controller evolution track* while minimizing collisions with static and dynamic obstacles and other vehicles. Static obstacles represented unexpected objects on the road and dynamic obstacles oncoming traffic and pedestrians.

Fitness Function: Controllers were awarded a fitness equalling the portion of the track’s length covered (via checkpoints) over 45 simulation (task trial) iterations:

$$\text{fitness}(x) = \frac{1}{\text{cars}} \sum_{i=0}^{\text{cars}} \left(\frac{cp_{\text{passed}}}{cp_{\text{total}}} * 0.9^{\text{coll}} \right) \quad (1)$$

Where, *cars* represents the number of vehicles, *cp_{passed}* the number of checkpoints vehicles successfully pass, *cp_{total}* the total number of check-points on that track, and *coll* the number of vehicle collisions (values lower than 0.9 resulted in slower evolution and often caused evolution to stagnate).

Novelty Search: Uses a *novelty metric* to determine a controller’s novelty [5], described by *sparseness* (equation 2), combining behavior characterization (vehicle speed) and a metric to compute a controller’s sparseness (novelty):

$$\text{Sparseness}(x) = \frac{1}{k} \sum_{i=0}^k \text{dist}(x, \mu_i) \quad (2)$$

Where, μ is the *i*th-nearest neighbor of *x* with respect to the novelty metric, and where the distance component in equation 2 uses the Euclidean distance.

Hybrid Search: Linearly combines novelty and fitness to create a weighted sum [4], where the score that controller *i* receives is defined by equation 3:

$$\text{score}(i) = \rho \cdot \overline{\text{fit}}(i) + (1 - \rho) \cdot \overline{\text{nov}}(i) \quad (3)$$

Where, $\rho = 0.5$, to equally combining (normalized) fitness and novelty for controller *i*:

$$\overline{\text{fit}}(i) = \frac{\text{fit}(i) - \text{fit}_{\text{min}}}{\text{fit}_{\text{max}} - \text{fit}_{\text{min}}}, \overline{\text{nov}}(i) = \frac{\text{nov}(i) - \text{nov}_{\text{min}}}{\text{nov}_{\text{max}} - \text{nov}_{\text{min}}} \quad (4)$$

3 EXPERIMENTS AND RESULTS

Two sets of experiments were conducted: *controller evolution* and *controller generalization test* experiments. The evolution experiments applied NEAT for vehicle controller evolution directed by *objective-based*, *novelty* or *hybrid search*, where average task performance was calculated over 20 runs. One experiment comprised

controller evolution coupled with one of these three evolutionary search methods, where each evolutionary run was 100 generations and each generation consisted of six simulation task trials that initialized three vehicles in random starting positions within a *starting area* at the start of the evolutionary evaluation track.

In the controller generalization test experiments, the highest task-performance evolved controller yielded after each of 20 evolutionary runs, was transferred to either *one, three* or *five* vehicles for non-evolutionary simulation test runs. Each generalization experiment was the 20 best controllers (evolved by a given NEAT search method) being executed (run in non-evolutionary task trial simulations) on three increasingly difficult variations of three test tracks.

Results indicate that controller evolution directed by hybrid search was significantly more expedient at evolving effective controllers compared to objective and novelty search (*Mann-Whitney U*, $p \leq 0.05$). An average (normalized) fitness of approximately 0.75 was reached by hybrid search after 20 generations, compared to 0.43 and 0.46 yielded by objective and novelty search, respectively. Results also indicate that all search methods yielded just above 60% task performance, where controller evolution directed by hybrid search significantly out-performed (*Mann-Whitney U*, $p \leq 0.05$) the objective and novelty search methods. However, there was no significant difference between objective and novelty search directed controller evolution. This result supports hybrid evolutionary search in this task and also previous work in multi-agent behavior evolution [2].

Generalization test results for the fittest controllers evolved by *objective, novelty* and *hybrid* search, indicated that vehicle controllers evolved with objective-based search were best able to generalize to the test tracks, whereas controllers evolved with hybrid search were least well suited to generalize to the test tracks. *Mann-Whitney U* ($p \leq 0.05$) tests indicated a significant difference between generalization test average task performance results of each search method, where the fittest controllers evolved by objective, hybrid and novelty search yielded an average task performance of 0.29, 0.20, and 0.24, respectively, in these generalization tests.

The efficacy of hybrid search is theorized to be a result of the thorough behavior space search. This exploratory capability of hybrid search is also supported by related work [2]. However, broad search space exploration enabled the discovery of minimally complex controllers achieving significantly higher average fitness, compared to objective and novelty search evolved controller behaviors. While this simple neural complexity was effective on the controller evolution track, controller generalization tests indicated that such simple controllers were ineffective across all test tracks.

The higher neural complexity of the fittest controllers evolved by objective and novelty search enabled sufficient behavioral functionality such that these controllers yielded significantly higher

average task performance across all test tracks. This result is similarly supported by related multi-agent behavior evolution work [2], demonstrating increased evolved controller complexity as detrimental to task performance on specific tasks (controller evolution track in this case), but generally beneficial to task performance across several tasks of varying difficulty (test tracks in this study).

Thus objective search was most effective for eliciting driving behaviors capable of generalizing to a broader set of related task environments (roads) of varying difficulty. This result does not support the general efficacy of hybrid-based controller evolution demonstrated in related multi-agent work [2], indicating that the coordinated driving task is not well suited to evolutionary search using behavioral diversity (novelty and hybrid search), but rather to objective-based search with a strictly defined fitness function.

An end goal of this research is develop evolutionary methods to automate autonomous vehicle controller design. Thus when such vehicles transit any given road or highway, an effective (safe) and efficient (expedient) coordinated-driving behavior emerges from vehicle interactions. Such automated controller design methods could potentially assist in designing fully distributed intelligent transportation systems, where autonomous vehicle manufacturers develop autonomous vehicle fleets that do not rely on costly centralized control systems for future roads or highways.

4 CONCLUSIONS

This study investigated controller automation for coordinated vehicle traffic on *autonomous vehicle only* roads. Results indicated that the fittest driving behaviors evolved by hybrid search did not generalize well to new test roads, compared the fittest driving behaviors evolved by objective-based and novelty search. The relatively poor performance of the fittest hybrid evolved controllers in these generalization tests indicates that while hybrid search was efficient and effective, the task constraints and variability of test tracks were not conducive to controllers evolved by hybrid search.

REFERENCES

- [1] T. Au, S. Zhang, and P. Stone. 2015. Autonomous Intersection Management for Semi-Autonomous Vehicles. In *Handbook of Transportation*, D. Teodorovic (Ed.). Taylor and Francis, London, UK, 88–104.
- [2] G. G. Nitschke and S. Didi. 2017. Evolutionary Policy Transfer and Search Methods for Boosting Behavior Quality: RoboCup Keep-Away Case Study. *Frontiers in Robotics and AI* 4, 62 (2017), 1–25.
- [3] A. Huang and G. Nitschke. 2017. Evolving Collective Driving Behaviors. In *Proceedings of the International Conference on Autonomous Agents and Multiagent Systems*. ACM, Sao Paulo, Brazil, 1573–1574.
- [4] C. Huang, D. Shorten, and G. Nitschke. 2015. Searching for novelty in pole balancing. In *Proceedings of the IEEE Congress on Evolutionary Computation*. IEEE, Sendai, Japan, 1792–1798.
- [5] J. Lehman and K. Stanley. 2011. Abandoning objectives: Evolution through the search for novelty alone. *Evolutionary computation* 19, 2 (2011), 189–223.
- [6] K. Stanley and R. Miikkulainen. 2002. Evolving neural networks through augmenting topologies. *Evolutionary Computation*. 10, 2 (2002), 99–127.